arXiv:1411.3434v1 [math.ST] 13 Nov 2014

On Approximations of the Beta Process in Latent Feature Models

Luai Al Labadi *& Mahmoud Zarepour†

November 14, 2014

ABSTRACT

The beta process has recently been widely used as a nonparametric prior for different models in machine learning, including latent feature models. In this paper, we prove the asymptotic consistency of the finite dimensional approximation of the beta process due to Paisley & Carin (2009). In addition, we derive an almost sure approximation of the beta process. This approximation provides a direct method to efficiently simulate the beta process. A simulated example, illustrating the work of the method and comparing its performance to several existing algorithms, is also included.

Keywords: Beta process; Ferguson and Klass representation; Finite dimensional approximation; Latent feature models; Simulation.

*Address for correspondence: Luai Al Labadi, Department of Statistical Sciences, University of Toronto, Toronto, Ontario M5S 3G3, Canada. E-mail: luai.allabadi@utoronto.ca.

†M. Zarepour, Department of Mathematics and Statistics, University of Ottawa, Ottawa, Ontario, K1N 6N5, Canada. E-mail: zarepour@uottawa.ca.

AMS Subject Classification: Primary 62F15, 62G20; secondary 60G51.

1 Introduction

The beta process was introduced by Hjort (1990) and later clarified by Kim (1999) in the field of Bayesian survival analysis as a class of prior processes for cumulative hazard function. Most recently, Kim et al. (2013) used the beta process to simulate the beta-Dirichlet process, a nonparametric prior for the cumulative intensity functions of a Markov process. Developments of the beta process in machine learning first appeared in the work of Thibaux & Jordan (2007), where it was shown through the application of document classification that the beta process could be used as a nonparametric prior in latent feature models. They also demonstrated that, when the beta process is marginalized out, one can obtain the Indian buffet process first defined in Griffiths & Ghahramani (2006). Since then, the beta process has been considered for many other applications in machine learning including factor analysis (Paisley & Carin, 2009), featural representations of multiple time series (Fox et al., 2009), Gene-expression analysis (Chen et al., 2010), linear regression (Chen et al., 2010), dictionary learning for image processing (Zhou et al., 2011), and image interpolation (Zhou et al., 2012).

Deriving a stick breaking representation for the beta process started by the work of Teh et al. (2007), who developed the stick breaking representation of the Indian buffet process. Recently, a stick-breaking construction of the full beta process was derived by Paisley et al. (2010). The derivation relied on a limiting process involving finite matrices, analogous to the limiting process used to derive the Indian buffet process. Broderick et al. (2012) demonstrated that the stick-breaking construction of the beta process can be directly obtained

from the characterization of the beta process as a Poisson process. A finite approximation of the beta process was suggested without proof by Paisley & Carin (2009). To this date, there is no mathematical proof for this approximation despite its use in several applications, including those previously mentioned. Providing a precise proof for the finite approximation of the beta process is the first goal of this paper.

Sampling from the beta process plays a central role in applications including latent feature models. For example, in factor analysis models (West, 2003; Paisley & Carin, 2009; Broderick et al., 2012), the data matrix is decomposed into the product of two matrices plus noise. The model takes the form:

$$X = Z\Phi + E,$$

where $X \in \mathbb{R}^{N\times P}$ is the data matrix and $E \in \mathbb{R}^{N\times P}$ is an error matrix. The matrix $\Phi \in \mathbb{R}^{K\times P}$ is a matrix of factors, and $Z \in \mathbb{R}^{N\times K}$ is a binary matrix of factor loadings. The dimension K is infinite, and thus the rows of Φ consist of an infinite collection of factors. The matrix Z is formed via a draw from a beta-Bernoulli process. First a sample from the beta process is drawn. Then applying this draw to a Bernoulli process yields an infinite binary vector of the matrix Z. The previous step is repeated to generate the matrix Z, where each successive draw of the Bernoulli process yields a further row of Z. In other words, the beta process is used to provide an infinite collection of coin-tossing probabilities. Tossing these coins corresponds to a draw from the Bernoulli process, yielding an infinite binary vector that is considered as a latent feature vector (Broderick et al., 2012). Deriving a simple, yet efficient, way to simulate the beta process is the second contribution for this paper.

The rest of the paper is organized as follows. In Section 2, we introduce the beta process and its conjugate process, the Bernoulli process. In Section 3, we prove the finite dimensional approximation of the beta process. In Section 4, an efficient and convenient method for

simulating the beta process is proposed. The approach is based on deriving a finite sum-representation which converges almost surely to the Ferguson & Klass representation (1972) of the beta process. An example illustrating the method and its performance to other existing approximations is presented in Section 5. Finally, our findings are briefly summarized in Section 6.

2 The Beta Process and the Bernoulli Process

The beta process and the Bernoulli process are examples of a general family of random measures known as *completely random measures*. Consider a space $\mathbb{X}$ with a $\sigma-$algebra $\mathscr{B}$ of subsets of $\mathbb{X}$. A random measure Φ is said to be completely random measure if for any finite collection $A_1, \ldots, A_n$ of disjoint members of $\mathscr{B}$, the random variables $\Phi(A_1), \ldots, \Phi(A_n)$ are independent. For more details about completely random measures, consult Kingman (1967). Let B_0 be a fixed continuous (non-atomic) finite measure on $(\mathbb{X}, \mathscr{B})$ and c be a positive number. Following Thibaux & Jordan (2007), the beta process B, written $B \sim BP(c, B_0)$, is a completely random measure with Lévy measure

$$\nu(d\omega, ds) = cs^{-1}(1-s)^{c-1} ds dB_0(\omega), \;\; 0 < s < 1, \omega \geq 0. \tag{2.1}$$

For any $S \in \mathbb{X}$, we have (Hjort, 1990; Thibaux & Jordan, 2007):

$$E[B(S)] = B_0(S) \quad \text{and} \quad Var[B(S)] = \frac{B_0(S)}{c+1}. \tag{2.2}$$

As in the Dirichlet process (Ferguson, 1973), c is called the *concentration parameter* and B_0 is called the *base measure*. Note that, in general, c can be a positive function of ω, but this

is not commonly used in latent feature models. The total mass of B_0, $\gamma := B_0(\mathbb{X})$, is called the *mass parameter*. A draw $B \sim \text{BP}(c, B_0)$ is described by:

$$B = \sum_{i=1}^{\infty} p_i \delta_{\omega_i}, \tag{2.3}$$

where $(p_1, \omega_1), (p_2, \omega_2), \ldots$ are the set of atoms in a realization of a nonhomogeneous Poisson process with mean measure ν. Here and throughout the present paper, δ_X denotes the Dirac measure at X, i.e. $\delta_X(A) = 1$ if $X \in A$ and 0 otherwise for a set $A \in \mathscr{B}$. As shown in (2.3), B is a discrete random measure (with probability 1). Note that, B is a finite measure since $E[\sum_{i=1}^{\infty} p_i] = E\left[B(\mathbb{X})\right] = B_0(\mathbb{X}) = \gamma$.

The stick-breaking representation of the beta process takes the form: (Paisley et al., 2010; Broderick et al., 2012)

$$B = \sum_{i=1}^{\infty} \sum_{j=1}^{C_i} V_{i,j}^{(i)} \prod_{l=1}^{i-1} (1 - V_{i,j}^{(l)}) \delta_{\omega_{i,j}}, \tag{2.4}$$

where $C_i \overset{i.i.d.}{\sim} Poisson(\gamma)$, $V_{i,j}^{(l)} \overset{i.i.d.}{\sim} beta(1, c)$, and $\omega_{i,j} \overset{i.i.d.}{\sim} B_0/\gamma$. When $C_i = 0$, then the corresponding sum is taken to be zero. The key difference between the stick-breaking representation of the Dirichlet process (Sethuraman, 1994) and that of the beta process is the weights (probabilities). The weights in the Dirichlet process depend on each other, while this is not the case for the beta process. Specifically, the weights that result from the stick-breaking representation of the Dirichlet process all come from a single stick (the unit interval). Thus, they add up to one. On the other hand, in the beta process, the weights all come from different unit intervals. So they need not add to one. However, as pointed out previously, their sum is finite almost surely (a.s.).

A connection between the beta process and the Poisson process was established in Paisley et al. (2012) and Broderick et al. (2012). They demonstrated that the beta process is a

Poisson process with the same mean measure (2.1). In particular, Paisley et al. (2012) showed the the stick-breaking construction defined in (2.4) is equivalent to

$$B \overset{d}{=} \sum_{j=1}^{C_1} V_{1,j}\delta_{\omega_{1,j}} + \sum_{i=2}^{\infty}\sum_{j=1}^{C_i} V_{i,j}e^{-T_{ij}}\delta_{\omega_{i,j}}, \tag{2.5}$$

where $V_{i,j} \overset{i.i.d.}{\sim} beta(1,c)$, $T_{ij} \overset{ind.}{\sim} gamma(i-1,c)$, C_i and $\omega_{i,j}$ are as defined in (2.4). In this paper, "$\overset{d}{=}$", "$\overset{d}{\to}$", "$\overset{v}{\to}$" and "$\overset{a.s.}{\to}$" denote equal in distribution, convergence in distribution, vague convergence and almost sure convergence, respectively. More details about convergence of random measures given in Appendix A. In addition, we use the same notation for the probability measure and its corresponding cumulative distribution function, i.e. $B(t) = B((-\infty, t])$ for $t \in \mathbb{X} = \mathbb{R}$. The inverse of a distribution function (or measure) B is defined by

$$B^{-1}(t) = \inf\{x : B(x) \geq t\}, \;\; 0 < t < 1.$$

A direct link to the Poisson process is the following representation:

$$B \overset{d}{=} \sum_{i=1}^{\infty}\sum_{j=1}^{C_i}\left(e^{-\Gamma_{i-1,j}/c} - e^{-\Gamma_{i,j}/c}\right)\delta_{\omega_{i,j}}, \tag{2.6}$$

where $\Gamma_{i,j} = E_{i,j}^{(1)} + \cdots + E_{i,j}^{(i)}$ and $(E_{i,j}^{(l)})_{1\leq l\leq i}$ are i.i.d. random variables with the exponential distribution with mean 1, independent of $\omega_{i,j}$. Here we have $\Gamma_{0,j} = 0$. The relationship (2.6) can be derived from the fact that

$$\begin{aligned} e^{-\Gamma_{i-1,j}/c} - e^{-\Gamma_{i,j}/c} &= e^{-E_{i,j}^{(1)}/c}\ldots e^{-E_{i,j}^{(i-1)}/c}\left(1 - e^{-E_{i,j}^{(i)}/c}\right) \\ &\overset{d}{=} (1-V_{i,j}^{(1)})\ldots(1-V_{i,j-1}^{(i-1)})V_{i,j}^{(i)}, \end{aligned}$$

since $e^{-E_{1,j}/c} \overset{d}{=} beta(c,1) \overset{d}{=} 1 - beta(1,c)$. See Iswaran & Zarepour (2000) for an analogous representation of the Dirichlet process. Expanding the summation in (2.6) for the first values of i gives:

$$\begin{aligned} B \overset{d}{=} & \sum_{j=1}^{C_1} \left(1 - e^{-\Gamma_{1,j}/c}\right) \delta_{\omega_{1,j}} \\ & + \sum_{j=1}^{C_2} \left(e^{-\Gamma_{1,j}/c} - e^{-\Gamma_{2,j}/c}\right) \delta_{\omega_{2,j}} \\ & + \sum_{j=1}^{C_3} \left(e^{-\Gamma_{2,j}/c} - e^{-\Gamma_{3,j}/c}\right) \delta_{\omega_{3,j}} + \cdots . \end{aligned}$$

Since, $\Gamma_{i,j}/i \overset{a.s}{\to} 1$ as $i \to \infty$ (by the strong law of large numbers), the weights becomes negligible for a large value of i. This makes representation (2.6) useful for simulation purposes through a truncation approach. See Zarepour & Al Labadi (2012) and Al Labadi & Zarepour (2013a) for further discussion about truncation procedures.

As pointed out earlier, the beta process is useful as a parameter for the Bernoulli process. The Bernoulli process can be defined, in general, for any base measure on $\mathbb{X}$. In our case, we consider the the base measure to be B, where $B \sim BP(c, B_0)$. Then a Bernoulli process Y with base measure B, written $Y \sim BeP(B)$, is a completely random measure

$$Y = \sum_i b_i \delta_{\omega_i},$$

where $b_i \overset{i.i.d.}{\sim} Bernoulli(p_i)$ for p_i given in (2.3). Observe that, since $E[Y|B] = B$ and B is a finite measure, the number of non-zero points in any realization of the Bernoulli process is finite.

The following theorem shows that the beta process is the conjugate prior for the Bernoulli

process. This conjugacy extends the conjugacy between the Bernoulli and beta distributions. For the proof of the theorem, see Thibaux & Jordan (2007) and Miller (2011).

Theorem 1. *Let $B \sim BP(c, B_0)$, and, for $i = 1, \ldots, m$, let $X_i|B \sim BeP(B)$ be m independent Bernoulli process draws from B. The posterior distribution of B after observing $X_1, \ldots, X_m$ is still a beta process:*

$$B_m^* = B|X_1, \ldots, X_m \sim BP\left(c^*, B_{0,m}^*\right),$$

with $c^* = c + m$ and

$$B_{0,m}^* = \frac{c}{c+m} B_0 + \frac{1}{c+m} \sum_{i=1}^{m} X_i.$$

3 Finite Dimensional Approximation of the Beta Process

In this section, we prove convergence of the finite approximation of the beta process, which was originally proposed by Paisley & Carin (2009) without a proof. As mentioned in the Introduction, this approximation plays a crucial role in several applications.

Theorem 2. *Consider a space $(\mathbb{R}, \mathcal{B})$, where $\mathbb{R}$ denotes the real line and $\mathcal{B}$ is the Borel $\sigma-$algeba of subsets of $\mathbb{R}$. Let B_0 be a finite continuous fixed measure on $(\mathbb{R}, \mathcal{B})$ with $B_0(\mathbb{R}) = \gamma$ and c be a positive number. For $n > \gamma$, define the process B_n as follows:*

$$B_n = \sum_{i=1}^{n} p_{i,n} \delta_{\omega_i},$$

$$p_{i,n} \overset{i.i.d.}{\sim} Beta\left(\frac{c\gamma}{n}, c\left(1 - \frac{\gamma}{n}\right)\right),$$

$$\omega_i \overset{i.i.d.}{\sim} B_0/\gamma,$$

$$(p_{i,n})_{1\leq i\leq n} \text{ and } (\omega_i)_{1\leq i\leq n} \text{ are independent.}$$

Then, as $n \to \infty$, $B_n \xrightarrow{d} B$, *where* $B \sim BP(c, B_0)$.

Proof. The proof is decomposed in several parts.

Part I: We apply Proposition 1 (Appendix A) to show that, as $n \to \infty$,

$$\begin{aligned} nP[p_{1,n} \in (x,1)] &= \frac{n\Gamma(c)}{\Gamma(c\gamma/n)\Gamma(c-c\gamma/n)} \int_x^1 s^{c\gamma/n-1}(1-s)^{c(1-\gamma/n)-1}ds \\ &\xrightarrow{v} \mu(x) = c\gamma \int_x^1 s^{-1}(1-s)^{c-1}ds. \end{aligned} \tag{3.1}$$

Observe that, for any $x > 0$, $\Gamma(x) = \Gamma(x+1)/x$. With $x = c\gamma/n$, we obtain $n/\Gamma(c\gamma/n) = c\gamma/\Gamma(c\gamma/n+1)$. Since $\Gamma(x)$ is a continuous function, as $n \to \infty$, we get

$$\frac{n}{\Gamma(c\gamma/n)} \to c\gamma$$

and

$$\frac{\Gamma(c)}{\Gamma(c-c\gamma/n)} \to 1.$$

It follows that, as $n \to \infty$,

$$\frac{n\Gamma(c)}{\Gamma(c\gamma/n)\Gamma(c-c\gamma/n)} \to c\gamma.$$

On the other hand, since $x < s < 1$, we have $s^{-1} < x^{-1}$ and $s^{c/n} < 1$. Thus, $s^{c/n-1} < x^{-1}$. Consequently, the integrand in (3.1) is dominated by $x^{-1}(1-s)^{c(1-1/n)-1}$, which is integrable for $x < s < 1$. Therefore, by the dominated convergence theorem, (3.1) holds. It follows

immediately from Proposition 1 that, as $n \to \infty$,

$$\xi_n = \sum_{i=1}^{n} \delta_{(p_{i,n},\omega_i)} \xrightarrow{d} \xi,$$

where ξ is a Poisson random measure with mean $d\mu \times dB_0/\gamma$ and ω_i is as defined in the statement of Theorem 2.

Part II: We show that

$$\xi = \sum_{i=1}^{\infty} \delta_{(\mu^{-1}(\Gamma_i),\omega_i)},$$

where $\Gamma_i = E_1 + \cdots + E_i$ and $(E_i)_{i\geq 1}$ is a sequence of i.i.d. random variables with the exponential distribution with mean 1. It is clear that $\sum_{i=1}^{\infty} \delta_{\Gamma_i}$ is a a PRM(λ), where λ is Lebesgue measure (Resnick, 2006, Example 5.1). By the special case of Proposition 3.9 of Resnick (1987), $\sum_{i=1}^{\infty} \delta_{(\Gamma_i,\omega_i)}$ is a PRM($\lambda \times B_0/\gamma$). Define $T : [0,\infty) \times (-\infty,\infty) \to [0,\infty) \times (-\infty,\infty)$ via $T(x,y) = (\mu^{-1}(x), y)$. If $t > 0$ and $a < b$, we have

$$\begin{aligned}
(\lambda \times B_0/\gamma) \circ T^{-1}([t,\infty) \times (a,b)) &= (\lambda \times B_0/\gamma)\left\{(x,y) : \mu^{-1}(x) \geq t \text{ and } a < y < b\right\} \\
&= (\lambda \times B_0/\gamma)\left\{(x,y) : x \leq \mu(t) \text{ and } a < y < b\right\} \\
&= \lambda([0,\mu(t)])\, B_0((a,b))/\gamma \\
&= \mu(t) B_0((a,b))/\gamma.
\end{aligned}$$

Thus, by Proposition 3.7 of Resnick (1987), $\xi = \sum_{i=1}^{\infty} \delta_{(\mu^{-1}(\Gamma_i),\omega_i)}$ is a PRM($d\mu \times dB_0/\gamma$). Therefore, as $n \to \infty$,

$$\xi_n = \sum_{i=1}^{n} \delta_{(p_{i,n},\omega_i)} \xrightarrow{d} \xi = \sum_{i=1}^{\infty} \delta_{(\mu^{-1}(\Gamma_i),\omega_i)}. \tag{3.2}$$

Part III: For $h > 0$, the map

$$T_h\left(\sum_{i=1}^{n}\delta_{(p_{i,n},\omega_i)}\right) = \sum_{i=1}^{n} p_{i,n} I\left(p_{i,n} > h\right)\delta_{\omega_i}$$

defined on the set of point processes is continuous with respect to vague topology for random measures (there are finite number of terms in the summation). Therefore, as $n \to \infty$ and for $h > 0$, applying T_h to (3.2), we obtain by the continuous mapping theorem (Resnick, 1987, p. 152)

$$B_{n,h} = \sum_{i=1}^{n} p_{i,n} I\left(p_{i,n} > h\right)\delta_{\omega_i} \xrightarrow{d} B_h = \sum_{i=1}^{\infty} \mu^{-1}(\Gamma_i) I\left(\mu^{-1}(\Gamma_i) > h\right)\delta_{\omega_i}.$$

Note that, as $h \to 0$, $B_h \xrightarrow{d} B$, where $B = \sum_{i=1}^{\infty} \mu^{-1}(\Gamma_i)\delta_{\omega_i}$. To complete the proof, by Theorem 3.2 of Billingsley (1999), it remains to show that for any Borel set A,

$$\begin{aligned} D_n^{h,\epsilon}(A) &= P\left\{|B_n(A) - B_{n,h}(A)| \geq \epsilon\right\} \\ &= P\left\{\left|\sum_{i=1}^{n} p_{i,n}\delta_{\omega_i}(A) - \sum_{i=1}^{n} p_{i,n} I\left(p_{i,n} \geq h\right)\delta_{\omega_i}(A)\right| \geq \epsilon\right\} \to 0, \end{aligned}$$

as $n \to \infty$ and $h \to 0$. We have

$$\begin{aligned} D_n^{h,\epsilon}(A) &= P\left\{\sum_{i=1}^{n} p_{i,n} I\left(p_{i,n} \leq h\right)\delta_{\omega_i}(A) \geq \epsilon\right\} \\ &\leq \epsilon^{-1} E\left[\sum_{i=1}^{n} p_{i,n} I\left(p_{i,n} \leq h\right)\delta_{\omega_i}(A)\right] \\ &\leq \epsilon^{-1}\sum_{i=1}^{n} E\left[p_{i,n} I\left(p_{i,n} \leq h\right)\right] E\left[\delta_{\omega_i}(A)\right] \\ &= \epsilon^{-1} n E\left[p_{1,n} I(p_{1,n} \leq h)\right] E\left[\delta_{\omega_1}(A)\right] \end{aligned}$$

$$= c\epsilon^{-1} E\left[\delta_{\omega_1}(A)\right] \int_0^h xnP\left\{p_{1,n} \in dx\right\}.$$

By (3.1), as $n \to \infty$, we get

$$D_n^{h,\epsilon}(A) \to c\epsilon^{-1} E\left[\delta_{\omega_1}(A)\right] \int_0^h (1-x)^{c-1} dx. \tag{3.3}$$

Observe that, the integral in (3.3) goes to zero as $h \downarrow 0$. Therefore, by Theorem 3.2 of Billingsley (1999), $B_n \xrightarrow{d} B$, where

$$B = \sum_i \mu^{-1}(\Gamma_i)\delta_{\omega_i} \tag{3.4}$$

is the Ferguson & Klass (1972) representation of the beta process. This completes the proof of the theorem. □

4 A New Algorithm to Generate the Beta Process

There are two general techniques to write a series representation for any Lévy process having no Gaussian component. The first one is due to Ferguson & Klass (1972). The second technique is due to Wolpert & Ickstadt (1998). Since we consider only the case when the concentration parameter is constant (i.e., $c(\omega) = c$ for all $\omega \geq 0$), the approach of Wolpert & Ickstadt (1998) and the approach of Ferguson & Klass (1972) are equivalent. See also Al Labadi & Zarepour (2013b, 2014a,b) for further discussion about the two approaches. Therefore, the algorithm of Wolpert & Ickstadt (1998) is not included here.

Let $B \sim BP(c, B_0)$ with a continuous B_0. Then the Ferguson & Klass (1972) representation of B takes the form given in (3.4). Since no closed form for the inverse of the Lévy measure (2.1) exists, working with (3.4) is relatively difficult in practice. The next theorem outlines a remedy to this problem, where an almost sure approximation to (3.4) is developed based on a similar result in Zarepour & Al Labadi (2012) for the Dirichlet process. Convergence of random measures is taken with respect to the vague topology on the space of point measures. Consult Appendix A for a background on convergence of random measures.

Theorem 3. *Let B_0 be a fixed continuous finite measure on $\mathbb{R}$ with $B_0(\mathbb{R}) = \gamma$. Let $(\omega_i)_{i\geq 1}$ be i.i.d. random variables with common distribution B_0/γ and $\Gamma_i = E_1 + \cdots + E_i$, where $(E_i)_{i\geq 1}$ are i.i.d. random variables with the exponential distribution with mean 1, independent of $(\omega_i)_{i\geq 1}$. For $n > \gamma$, define*

$$\mu_n(x) = \frac{\Gamma(c)}{\Gamma(\gamma c/n)\Gamma(c - c\gamma/n)} \int_x^1 s^{c\gamma/n-1} (1-s)^{c(1-\gamma/n)-1} \, ds.$$

Then, as $n \to \infty$,

$$B_n = \sum_{i=1}^{n} \mu_n^{-1}\left(\frac{\Gamma_i}{\Gamma_{n+1}}\right)\delta_{\omega_i} \overset{a.s.}{\to} B = \sum_{i=1}^{\infty} \mu^{-1}(\Gamma_i)\,\delta_{\omega_i}.$$

Proof. It follows from the proof of Theorem 2 that, for any $x > 0$,

$$n\mu_n(x) \to \mu(x). \tag{4.1}$$

Notice that, the left hand side of (4.1) is a sequence of a continuous monotone functions converging to a monotone function for every $x > 0$. This is equivalent to the convergence of their inverse function to the inverse function of the right hand side (Resnick, 1987, Proposi-

tion 0.1). Thus,

$$\mu_n^{-1}\left(\frac{x}{n}\right) \to \mu^{-1}(x). \tag{4.2}$$

Now, taking $x = \Gamma_i$ in (4.2) and the fact that $\Gamma_{n+1}/n \to 1$ as $n \to \infty$ (by the strong law of large numbers) we get

$$\mu_n^{-1}\left(\frac{\Gamma_i}{\Gamma_{n+1}}\right) \to \mu^{-1}(\Gamma_i). \tag{4.3}$$

To prove Theorem 2, by Lemma 1 of Al Labadi and Zarepour (2013b), we show that, for all k fixed,

$$\sum_{i=1}^{k} \mu_n^{-1}\left(\frac{\Gamma_i}{\Gamma_{n+1}}\right)\delta_{\omega_{(i)}} \overset{a.s.}{\to} \sum_{i=1}^{k} \mu^{-1}(\Gamma_i)\,\delta_{\omega_{(i)}}$$

as $n \to \infty$, where $\omega_{(1)} \leq \cdots \leq \omega_{(n)}$ represent the corresponding order statistics of $\omega_1, \ldots, \omega_n$. This directly follows by (4.3). □

The next algorithm is used to generate samples from an approximation of the beta process with parameters c and B_0, where B_0 continuous. The steps of the algorithm are:

(1) Fix a relatively large positive integer n.

(2) Generate $\omega_i \overset{\text{i.i.d.}}{\sim} B_0/\gamma$ for $i = 1, \ldots, n$.

(3) Generate $E_i \overset{\text{i.i.d.}}{\sim} exponential(1)$ for $i = 1, \ldots, n+1$ such that $(E_i)_{1\leq i\leq n+1}$ $(\omega_i)_{1\leq i\leq n}$ are independent where $\Gamma_i = E_1 + \cdots + E_i$.

(4) For $i = 1, \ldots, n$, compute $\left(\mu_n^{-1}(\Gamma_i/(\Gamma_{n+1}))\right)$, which is simply the quantile function of the $beta\,(c\gamma/n, c(1-\gamma/n))$ distribution evaluated at $1 - \Gamma_i/\Gamma_{n+1}$.

(5) Set B_n as in Theorem 3.

5 Empirical Results: Comparison with Other Methods

Several algorithms to sample from the beta process exist in the literature. In this section, we compare the new approximation of the beta process with the finite dimensional approximation of the beta process (Paisley & Carin,2009), Damien et al.'s (1995) algorithm, Lee & Kim's (2004) algorithm, and Lee's (2007) algorithm. Summary of these algorithms are given in Appendix B. We would like to highlight here that the new approximation converges almost surely to the Ferguson & Klass (1972) representation of the beta process, while the other four algorithms converge only in distribution. Almost sure convergence is not only stronger than convergence in distribution but also provides a proper path by path comparison between the suggested approximation and the limit. Thus, almost sure convergence is preferred over convergence in distribution.

In order to make comparisons among the algorithms, we use equivalent settings for the parameters characterizing these algorithms (see Table 1). We refer the reader to the original papers for the details of the algorithms. We consider the beta process with $c = 2$ and $B_0(x) = x$ (i.e., a uniform distribution on $[0, 1]$). We compute the absolute maximum difference between an approximate sample mean and the exact mean. See also Lee & Kim (2004) and Lee (2007) for similar comparisons. The exact mean is x; see (2.2). We refer to this statistic by the maximum mean error. Specifically,

$$\text{maximum mean error} = \max_x \left| E\left[B_n(x)\right] - E\left[B(x)\right] \right| \;\; = \;\; \max_x \left| E\left[B_n(x)\right] - x \right|,$$

where $x = 0.1, 0.2, \ldots, 0.9, 1.0$, B_n is an approximation of the beta process, and $B \sim BP(c = 2, B_0(x) = x)$. Note that $E\left[B_n(x)\right]$ is approximated by obtaining the mean at $x = 0.1, 0.2, \ldots, 0.9, 1.0$ of 3000 i.i.d. sample paths from the approximated process B_n.

Similarly, we compute the maximum standard deviation error between an approximate sample standard deviation (s.d.) and the exact standard deviation. The exact standard deviation is $\sqrt{x/3}$; see (2.2). Thus,

$$\text{maximum s.d. error} = \max_x |s.d\,[B_n(x)] - s.d.\,[B(x)]| \;\; = \;\; \max_x \left| s.d.\,[B_n(x)] - \sqrt{x/3} \right|.$$

Table 1 depicts values of the maximum mean error, the maximum standard deviation error, and the corresponding computational time. The computational time is computed by applying the code "System.Time" available in R. As seen in Table 1, the new algorithm has the smallest mean and standard deviation errors. In addition, it has a reasonable computation time.

Table 1: This table reports the maximum mean error, the maximum standard deviation error, and the corresponding computation time. Here, DSL, KL, and PC stand for Damien et al.'s algorithm (1995), Lee & Kim's algorithm (2004), and Paisley & Carin's (2009) algorithm, respectively.

Algorithm	Parameters	max. mean error	max. s.d. error	Time
DSL	$m = n = 200$	0.0110	0.0156	100.70
LK	$\epsilon = 0.01$	0.0128	0.0107	0.51
Lee	$n = 200, \epsilon = 0.05$	0.0162	0.0392	1.73
PC	$n = 200$	0.0113	0.1280	0.42
New	$n = 200$	0.0087	0.0061	26.38

6 Conclusions

In this paper, we have proved the finite dimensional approximation of the beta process (Paisley & Carin, 2009). This approximation has been used in several machine learning models. We also have derived an almost surely approximation of the beta process. This new approximation provides a simple, yet efficient, way to simulate the beta process.

REFERENCES

[1] Al Labadi, L. and Zarepour, M. (2013a). On asymptotic properties and almost sure approximation of the normalized inverse-Gaussian process. *Bayesian Anal.* **8**, 3, 553-568.

[2] Al Labadi, L. and Zarepour, M. (2013b). A Bayesian nonparametric goodness of fit test for right censored data based on approximate samples from the beta-Stacy process. *Canad. J. Statist.* **41**, 3, 466-487.

[3] Al Labadi, L., and Zarepour, M. (2014a). On Simulations from the Two-Parameter Poisson-Dirichlet Process and the Normalized Inverse-Gaussian Process. *Sankhya A* **76**, 158-176.

[4] Al Labadi, L. and Zarepour, M. (2014b). Goodness of fit tests based on the distance between the Dirichlet process and its base measure. *J. Nonparametr. Stat.* **26**, 341-357.

[5] Applebaum, D. (2004). *Lévy processes and Stochastic Calculus*. Cambridge University Press, Cambridge.

[6] Billingsley, P. (1999). *Convergence of Probability Measures*, third edition. John Wiley & Sons, Inc.

[7] Damien, P., Laud, P. and Smith, A. F. M. (1995). Approximate random variate generation from infinitely divisible distributions with applications to Bayesian inference. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **57**, 547-563.

[8] Broderick, T, Jordan, M. I., and Pitman, J. (2012). Beta processes, stick-breaking, and power laws. *Bayesian Anal.* **7**, 439-476.

[9] Chen, B., Paisley, J. and Carin, L. (2010). Sparse linear regression with beta process priors. International Conference on Acoustics, Speech, and Signal Processing, Dallas, TX.

[10] Chen, B., Chen, M., Paisley, J., Zaas, A., Woods, C., Ginsburg, G. S., Hero, A., Lucas, J., Dunson, D. and Carin, L. (2010). Bayesian inference of the number of factors in gene-expression analysis: Application to human virus challenge studies. *BMC Bioinformatics*, **11**:552.

[11] Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1**, 209-230.

[12] Ferguson, T. S. and Klass, M. J. (1972). A representation of independent increment processes without Gaussian components. *Annals of Mathematical Statistics* **1**, 209-230.

[13] Fox, E., Sudderth, E. , Jordan, M., and Willsky, A. (2009). Sharing features among dynamic systems with beta processes. *Adv. Neural Inf. Process. Syst.* **22**, 549-557.

[14] Griffiths, T. and Ghahramani, Z. (2006). Infinite latent feature models and the Indian buffet process. *Adv. Neural Inf. Process. Syst.* **18**, 475-482.

[15] Hjort, N. L. (1990). Nonparametric Bayes estimators based on Beta processes in models for life history data. *Ann. Statist.* **18**, 1259-1294.

[16] Ishwaran, H. , James, L.F., and Zarepour, M. (2009). An alternative to m out of n bootstrap. *J. Stat. Plan. Infer.*, **39**, 788-801.

[17] Ishwaran, H. and Zarepour, M. (2002). Exact and approximate sum representations for the Dirichlet process. *Canad. J. Statist.* **30**, 269-283.

[18] Kim, Y. (1999). Nonparametric Bayesian estimators for counting processes. *Ann. Statist.* **27**, 562-588.

[19] Kim, Y., James, L. and Weibbach, R. (2012). Bayesian analysis of multistate event history data: Beta-dirichlet process prior. *Biometrika* **99**, 127-140.

[20] Kingman, J. F. C. (1967). Completely random measures. *Pacific J. Math.* **21**, 59-78.

[21] Lee, J. (2007). Sampling methods for neutral to the right processes. *J. Comput. Graph. Statist.* **16**, 656-671.

[22] Lee, J. and Kim, Y. (2004). A new algorithm to generate beta processes. *Comput. Statist. Data Anal.* **25**, 401-405.

[23] Miller, K. T. (2011). Bayesian Nonparametric Latent Feature Models. PhD thesis, University of California, Berkeley.

[24] Paisley, J., Blei, D. and Jordan, M. (2012). Stick-breaking beta processes and the Poisson process. In Proceedings of the International Conference on Artifcial Intelligence and Statistics (AISTATS)

[25] Paisley, J. and Carin, L. (2009). Nonparametric factor analysis with beta process priors. In Proceedings of the International Conference on Machine learning (ICML).

[26] Paisley, J., Zaas, A., Woods, C. W., Ginsburg, G. S. and Carin, L. (2010). A stick-breaking construction of the beta process. In Proceedings of the International Conference on Machine learning (ICML).

[27] Resnick, S. I. (1987). *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag, New York.

[28] Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statist. Sinica* **4**, 639-650.

[29] Teh, Y. W., Görür, D. and Ghahramani, Z. (2007). Stick-breaking construction for the Indian buffet process. In Proceedings of the International Conference on Artifcial Intelligence and Statistics (AISTATS).

[30] Thibaux, R. and Jordan, M. I. (2007). Hierarchical beta processes and the Indian buffet process. In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS).

[31] van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes with Applications to Statistics*. Springer, New York.

[32] West, M. (2003). Bayesian factor regression models in the"large p, small n" paradigm. *Bayesian Anal.* **7**, 723-732.

[33] Wolpert, R. L. and Ickstadt, K. (1998). Simulation of Lévy random fields. In *Practical Nonparametric and Semiparametric Bayesian Statistics* (Edited by D. Day, P. Nuller and D. Sinha), 237-242. Springer, New York.

[34] Zarepour, M. and Al Labadi, L. (2012). On a rapid simulation of the Dirichlet process. *Statist. Probab. Lett.* **82**, 916-924.

[35] Zhou, M., Yang, H., Sapiro, G., Dunson, D. and Carin, L. (2011). Dependent hierarchical beta process for image interpolation and denoising. In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS).

[36] Zhou, M., Chen, H., Paisley, J., Ren, L., Li, L., Xing, Z., Dunson, D., Sapiro, G. and Carin, L. (2012). Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images. *IEEE Trans. Image Process* **21**,130-144.

A Vague Convergence

The main objective of this appendix is to give a brief introduction about convergence of random measures. Let $(\mathbb{E}, \mathscr{E})$ be a state space as before. Let $C_K^+(\mathbb{E})$ be the set of continuous functions $f : \mathbb{E} \to [0, \infty)$ with compact support. A measure μ is called Radon if $\mu(K) < \infty$ for any compact set K in $\mathbb{E}$. Let $M_+(\mathbb{E})$ be the space of Radon measures in $\mathbb{E}$. Let $\mathscr{M}_+(\mathbb{E})$ be the smallest $\sigma-$algebra of subsets of $M_+(\mathbb{E})$ making the maps $\mu \to \mu(f) = \int f(x)d\mu(x)$ from $M_+(\mathbb{E})$ to $\mathbb{R}$ measurable for all functions $f \in C_K^+(\mathbb{E})$. Note that, $\mathscr{M}_+(\mathbb{E})$ is the Borel

$\sigma-$algebra generated by the topology of vague convergence. If $\mu_n, \mu \in M_+(\mathbb{E})$, we say that $(\mu_n)_n$ converges vaguely to μ (and we write $\mu_n \xrightarrow{v} \mu$) if $\mu_n(f) \xrightarrow{v} \mu(f)$ for any $f \in C_K^+(\mathbb{E})$.

A *random measure* on $\mathbb{E}$ is any measurable map ξ defined on a probability space $(\Omega, \mathscr{A}, P)$ with values in $(M_+(\mathbb{E}), \mathscr{M}_+(\mathbb{E}))$. If ξ_n, ξ are random measures on $\mathbb{E}$, we say that $(\xi_n)_n$ converges in distribution to ξ (and we write $\xi_n \xrightarrow{d} \xi$) if $\{P \circ \xi_n^{-1}\}_n$ converges weakly to $P \circ \xi^{-1}$. By Theorem 4.2 of Kallenberg (1983), $\xi_n \xrightarrow{d} \xi$ if and only if $\xi_n(f) \to \xi(f)$, i.e.

$$\int_E f(x)\xi_n(dx) \to \int_E f(x)\xi(dx), \ \forall f \in C_K^+(\mathbb{E}).$$

We say that $(\xi_n)_n$ converges vaguely almost surely to ξ (and write $\xi_n \xrightarrow{a.s.} \xi$) if there exists a set $\widetilde{\Omega} \in \mathscr{A}$ with $P(\widetilde{\Omega}) = 1$ such that $\forall \omega \in \widetilde{\Omega}$, $\xi_n(\omega, \cdot) \xrightarrow{v} \xi(\omega, \cdot)$, i.e.

$$\int_E f(x)\xi_n(\omega, dx) \to \int_E f(x)\xi(\omega, dx), \ \forall f \in C_K^+(\mathbb{E}).$$

The space $M_+(\mathbb{E})$ endowed with the vague topology is a complete separable metric space (Resnick, 1987, Proposition 3.17). See also Kallenberg (1983).

The next proposition is fundamental in studying the weak convergence of Poisson processes. It gives necessary and sufficient conditions for empirical measures to converge to a Poisson random measure. The proof follows by mimicing the proof of Proposition 3.21 of Resnick (1987) with $j \geq 1$ and j/n are replaced by $1 \leq j \leq n$ and ω_j, respectively. See also Ishwaran, James, and Zarepour (2009) for an analogous result.

Proposition 1. *Suppose for each $n \geq 1$, we have $(X_{j,n})_{1 \leq j \leq n}$ are i.i.d. random variable random elements of $(\mathbb{E}, \mathscr{E})$ and μ is a Radon measure on $(\mathbb{E}, \mathscr{E})$. Let $\omega_i \overset{i.i.d.}{\sim} F$ such that $(\omega_j)_{1 \leq j \leq n}$ and $(X_{j,n})_{1 \leq j \leq n}$ are independent. Define $\xi_n = \sum_{j=1}^n \delta_{(\omega_j, X_{j,n})}$ and suppose ξ is a*

Poisson random measure on $\mathbb{R} \times \mathbb{E}$ *with mean* $dF \times d\mu$*. Then*

$$\xi_n \xrightarrow{d} \xi \text{ in } M_p(\mathbb{R} \times \mathbb{E}) \quad \text{if and only if} \quad nP[X_{1,n} \in \cdot] \xrightarrow{v} \mu \quad \text{on } \mathbb{E}.$$

B Other Sampling Algorithms

Several algorithms are suggested to sample from the beta process $B \sim BP(c, B_0)$ with a continuous B_0. We consider the algorithm of Damien, Laud, and Smith (1995), the algorithm of Lee and Kim (2004) and the algorithm Lee (2007). Below is a brief discussion of these algorithms. We refer the reader to the original papers for more details.

• **Damien-Laud-Smith Algorithm:** Using the fact that the distributions of the increments of a nondecreasing Lévy process are infinitely divisible, Damien, Laud, and Smith (1995) derived an algorithm to generate approximations for infinitely divisible random variables and used it to generate the beta process. Let p_i denotes the increment of the process B in the interval $\Delta_i = (\omega_{i-1}, \omega_i]$, i.e. $p_i = B(\omega_i) - B(\omega_{i-1})$. The steps of the Damien-Laud-Smith algorithm for simulating an approximation for the jump p_i are:

(1) Fix a relatively large positive integer n.

(2) Generate independent values z_{ij} from the probability density function $dB_0(t)/\, B_0(\Delta_i)$, for $j = 1, \ldots, n$.

(3) Generate $x_{ij} \sim beta(1, c)$, for $j = 1, \ldots, n$.

(4) Generate y_{ij}: $y_{ij}|x_{ij} \sim Poisson(\lambda_i n^{-1} x_{ij}^{-1})$, for $j = 1, \ldots, n$, where $\lambda_i = B_0(\omega_i) - B_0(\omega_{i-1})$.

(5) Set $p_{i,n} = \sum_{j=1}^{n} x_{ij} y_{ij}$. For large n, $p_{i,n}$ is an approximation of p_i.

Damien, Laud, and Smith (1995) showed that $p_{i,n} \xrightarrow{d} p_i$, as $n \to \infty$. That is, $p_{i,n}$ is an approximate sample from the ith increment of B. Note that, the Damien-Laud-Smith algorithm generates only the increments of the process and not the entire process. To obtain the whole process, we set

$$B_{m,n} = \sum_{i=1}^{m} p_{i,n}\delta_{\omega_i}.$$

For large m and n, $B_{m,n}$ is an approximation of B.

• **Lee-Kim Algorithm:** The Kim and Lee algorithm for the beta process with parameters c and B_0 with B_0 continuous can be described as follows. First the Lévy measure ν of the beta process given by (2.1) is approximated by:

$$\nu_\epsilon(d\omega, ds) = \frac{c}{\epsilon} b(s : \epsilon, c) dB_0(\omega) ds,$$

where

$$b(x : a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1}, \quad \text{for } 0 < x < 1, a > 0, b > 0. \tag{B.1}$$

The steps of the Lee-Kim algorithm for the beta process B are:

(1) Fix a relatively small positive number ϵ.

(2) Generate the total number of jumps $n \sim Poisson\,(c\gamma/\epsilon)$.

(3) Generate i.i.d. random variables $\omega_1, \ldots, \omega_n$ from the probability density function dB_0/γ.

(4) Let $\omega_{(1)} \leq \ldots \leq \omega_{(n)}$ be the corresponding order statistics of $\omega_1, \ldots, \omega_n$.

(5) Generate the jump sizes $p_1, \ldots, p_n : p_i|\omega_{(i)} \sim Beta(\epsilon, c)$.

(6) Set $B_\epsilon = \sum_{i=1}^{n} p_i \delta_{\omega_{(i)}}$.

Lee and Kim (2004) showed that $B_\epsilon \xrightarrow{d} B$, as $\epsilon \to 0$.

• **Lee Algorithm:** The steps of the Lee algorithm are:

(1) Fix a relatively large positive integer n.

(2) Generate i.i.d. random variables $\omega_1, \ldots, \omega_n$ from the probability density function dB_0/γ.

(3) For $i = 1, \ldots, n$, generate $x_i \sim b(s : \epsilon, c)$, where $b(s : \epsilon, c)$ is defined in (B.1)

(4) For $i = 1, \ldots, n$, generate $y_i \sim Poisson\left(\gamma b(x_i : 1, c)/n x_i b(x_i : \epsilon, c)\right)$.

(5) Set $B_n = \sum_{i=1}^{n} x_i y_i \delta_{\omega_i}$.

Lee (2007) proved that, as $n \to \infty$, $B_n \xrightarrow{d} B$.